Neutron Beam Testing of High Performance Computing Hardware

Sarah E. Michalak, Andrew J. DuBois, Curtis B. Storlie, Heather M. Quinn, William N. Rust, David H. DuBois, David G. Modl, Andrea Manuzzato, and Sean P. Blanchard

***Abstract*—Microprocessor-based systems are the most common design for high-performance computing (HPC) platforms. In these systems, several thousands of microprocessors can participate in a single calculation that could take weeks or months to complete. When used in this manner, a fault in any of the microprocessors could cause the computation to crash or cause silent data corruption (SDC), i.e. computationally incorrect results. In recent years, neutron-induced failures in HPC hardware have been observed, and researchers have started to study how neutron radiation affects microprocessor-based scientific computations. This paper presents results from an accelerated neutron test focusing on two microprocessors used in Roadrunner, the first Petaflop system.**

***Index Terms*—Soft error, single event effect, silent data corruption, neutron beam testing, cross-section.**

I. Introduction

Microprocessors are a crucial aspect of many different types of computation, including standard desktop computers, high-performance computing (HPC) platforms, avionics and spacecrafts. While the most common microprocessors used for space applications are radiation-hardened and have low SEU rates [1]–[5], commercial-off-the-shelf (COTS) microprocessors are more common in terrestrial and avionics computing. Radiation-induced failure modes in COTS microprocessors include single-event upsets (SEUs) in the caches, register files, pipeline registers, and memory; single-event transients (SETs) in functional units; and single-event functional interrupts (SEFIs) in control logic. Although the terrestrial radiation environment has a much lower flux than in space environments [6], neutron-induced faults are a concern for large terrestrial systems and neutrons have been implicated in crashes and silent data corruption (SDC) in different systems [7]–[9].

Because SEUs are increasingly noticeable in terrestrial applications, COTS microprocessor designers and system designers often include some protection from SEUs. These protections include error-correcting codes (ECC) and bit interleaving in caches. As many COTS microprocessors have multiple levels of caches, some manufacturers apply a graded approach to error mitigation, with the larger caches that are overwritten less frequently ECC protected and smaller caches that are frequently overwritten more weakly able to detect errors. In HPC platforms, software-level protections, such as checkpoint/restart, are also implemented. In checkpoint/restart, the calculation's state is periodically saved to hard disk so that a calculation can be restarted from the previous state if necessary.

While these protections are useful, not all radiation-induced failure modes are completely suppressed. HPC platforms used for scientific computation are particularly sensitive to these problems due to the size of the system and availability requirements of the applications run on them. Commonly these systems distribute a single calculation across thousands of microprocessors, which can spend weeks to months processing the calculation. The large system size can magnify tiny sensitivities to radiation in the system. Often if a single microprocessor crashes the calculation will be stopped, the previous checkpoint loaded, and the calculation restarted. For

Corresponding author S. E. Michalak is with the Statistical Sciences Group, Los Alamos National Laboratory, Los Alamos, NM 87544 USA (505-667-2625; fax: 505-667-4470; e-mail: michalak@lanl.gov). A. J. DuBois, D. H. DuBois, D. G. Modl, and S. P. Blanchard are with the System Integration Group, Los Alamos National Laboratory (e-mails: ajd@lanl.gov, dhd@lanl.gov, digem@lanl.goc, seanb@lanl.gov). C. B. Storlie and W. N. Rust are with the Statistical Sciences Group, Los Alamos National Laboratory (e-mails: storlie@lanl.gov, wnr@lanl.gov). H. M. Quinn is with the Space Data Systems Group, Los Alamos National Laboratory (e-mail: hquinn@lanl.gov). A. Manuzzato was with Universita degli Studi di Padova Italy (e-mail: andrea.manuzzato@ieee.org).

978-1-4577-1283-8/11/$26.00 ©2011 IEEE

these situations, every time the calculation is restarted from a checkpoint increases the application runtime, as all of the runtime between the checkpoint and the restart is lost. Furthermore, SDC can be difficult to detect. Therefore, in these platforms, neutron-induced errors are of concern since (1) system crashes affect application runtimes and (2) SDC in scientific applications may lead to incorrect scientific conclusions. The ability to infer field experience from accelerated testing data such as that resulting from neutron beam testing is important.

This paper presents results from beam testing hardware identical to that used in Roadrunner [10], the first Petaflop system [11], while running different applications including some used for scientific research. Further results are available in [12], [13]. Section II discusses related work on other microprocessor radiation studies. Section III presents the test setup, with the results detailed in Section IV. Section V concludes this work.

II. Related Work

There is more than a decade's worth of static test data on microprocessors [14]–[17]. A number of recent publications have studied more modern microprocessors with reduced feature sizes and multiple processing cores [18], [19]. While static testing is often the basis for error rate calculations, it can be difficult to translate these errors into dynamic error rates for real-world systems. Determining the overall effect of radiation on microprocessors is not simple, as faults in the systems can remain dormant for several thousands of clock cycles before triggering an error. In addition, the operating system and the software can create noise in the system, making it difficult to determine the cause of system crashes. In [20] results from [14] are used to indicate that the proton cross-section for the Pentium II and MMX microprocessors was two to three orders of magnitude larger when tested with Windows operating system than without.

With the ever decreasing transistor size in microprocessors, much has been done to improve radiation reliability. When compared to the state-of-the-art, there are two main methods to increase microprocessor reliability. The first method focuses on improving the radiation reliability of the microprocessor. The second focuses on improving radiation reliability through algorithm-based fault-tolerance.

While more reliable hardware would improve software reliability, many reliability improvements at the hardware level are prohibitively expensive to produce or cannot provide the performance of current non-radiation hardened devices. BAE Systems produces the RAD750 line of radiation-hardened microprocessors, which are based on the PowerPC 750 architecture. These microprocessors run at 200 MHz, and are advertised as achieving more than 420 Dhrystone MIPS [2]. An earlier, 133 MHz version of the microprocessor consumed 5W of power [1]. There are also a number of new radiation-hardened microprocessors that will be released soon, such as the radiation-hardened by design Maestro chip [4] or the radiation-hardened by process Clearspeed [5]. As all of these microprocessors are much slower than COTS microprocessors and performance is important for large-scale computations, using radiation-hardened microprocessors for HPC platforms is unlikely.

The second method that may be implemented is algorithm-based fault-tolerance, but it relies on specialized knowledge by the programmer or ad hoc optimization of the code by hand [21], [22]. These methods usually decrease performance, which is paramount in HPC systems, and are typically not used in HPC systems.

III. Test Setup and Experimental Protocol

Hardware from Roadrunner was tested at Los Alamos National Laboratory's (LANL) Los Alamos Neutron Science Center (LANSCE) Irradiation of Chips and Electronics (ICE) House in October 2009 to investigate the neutron susceptibility of the two microprocessors used in Roadrunner along with the hardware in their respective beampaths. Both microprocessors, the IBM PowerXCell 8i (Cell) and the AMD Opteron 2210 HE, have been commercially available. The test setup included running different applications, including some used for scientific computation.

The Cell is a 65nm SOI microprocessor with 1 PowerPC processor element (PPE) that controls 8 synergistic processor elements (SPE). See [23, p. 5] for a diagram of the Cell architecture. The 3.2 GHz PPE includes a PowerPC processor unit (PPU) that is based on the PowerPC architecture, a parity-protected 32 KB L1 data cache, a parity-protected 32 KB L1 instruction cache, and a 512 KB L2 cache with ECC on data and parity on directory tags (which is recoverable using redundant directories). Each 3.2 GHz SPE includes a synergistic processor unit (SPU) and an ECC-protected 256 KB dedicated non-caching local store. The IBM QS22 blade (system) housing the Cells during the testing includes 8 GB of ECC double data rate 2 (DDR2) dynamic random access memory (DRAM).

The Opteron 2210 HE is a 1.8 GHz 90nm Silicon-on-Insulator (SOI) dual-core microprocessor. See [24, p. 2] for a diagram of the Opteron 2210 EE microprocessor, which has a design similar to the Opteron 2210 HE

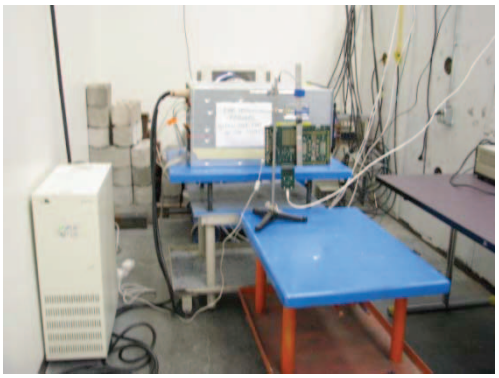

Fig. 1. The test setup, showing BC-H that housed the Triblades and the down-beam Virtex-II; photo from [25].

Fig. 2. Triblade 3 in BC-H; photo from [25].

tested at LANSCE. Each Opteron core has an ECC-protected 64 KB L1 data cache, a parity-protected 64 KB L1 instruction cache, and an ECC protected 1MB L2 cache. The IBM LS21 blade (system) housing the Opteron 2210 HE for the testing includes 16 GB of ECC DDR2 DRAM.

A. Test Setup

There are two aspects of the test setup: the hardware test setup and the software test setup. Both aspects were specifically designed to mimic how the devices under test operate in the Roadrunner supercomputer. This section presents both the hardware test setup and the software test setup.

1) Hardware Test Setup: Due to their extreme size, most HPC platforms need an efficient power, cooling and network design, as the entire system might span thousands of square feet. To this end, most platforms are designed to be housed in server racks, with each rack housing multiple chassis. In a blade-based platform such as Roadrunner, each chassis will house multiple blades. In Roadrunner, the rack provides physical structure, the chassis provides a common interface to power, network, and cooling, and compute blades provide the compute infrastructure.

Figs. 1 and 2 show the hardware test setup. The hardware tested included four Triblade [10] blade servers and a BladeCenter-H (BC-H Type 8852) [26] chassis. The BC-H is designed to house up to three Triblades, where each Triblade includes one IBM LS21 blade, two IBM QS22 blades, and an expansion blade to manage data traffic. The LS21 blade has two dual-core Opteron 2210 HE microprocessors. The QS22 blades (QS22a and QS22b) each have two Cell microprocessors. During testing, the Triblade under test was housed in the BC-H as it would be in the field. The BC-H was oriented so that with a single Triblade under test, the beam first entered the QS22b, followed by the QS22a, the expansion blade, and the LS21 respectively.

The test also required extra hardware for system control and radiation measurements. In many HPC platforms, a front-end node is used to manage the back-end compute nodes or blades, and likewise for this testing a front-end node was necessary to control the system under test. Specifically, it was used to boot the Triblades, start applications on the system under test and to monitor its health, all of which were performed manually by the experiment personnel. An IBM eServer X Series 336 provided this capacity. This server was placed in the user facility so that it would be protected from radiation.

Since the hardware setup included more physical matter (chassis, metal enclosures, large heatsinks) than usual, two Xilinx Virtex-II [27] test fixtures [28] were included, one placed upbeam of the BC-H and the other placed downbeam of the BC-H, for calculating corrected neutron fluence exposures. Normally, the decrease in radiation would be based on the distance from the beam source. Without the BC-H chassis and the Triblades, the expected decrease in radiation from the beam source to the back Virtex-II would be 20%. With the BC-H and one Triblade between the two Virtex-II devices, the back Virtex-II device indicated an average decrease in radiation of 69% from the beam source.

Because the decrease in radiation is larger than that expected given the distance from the beam source, the fluence was adjusted based on both distance and attenuation through matter. The decrease in radiation based on distance was calculated as usual. The decrease in radiation due to attenuation through matter was based on exponential decay. So that the attenuation effect would be tractable, the beam was assumed to decay in a uniform exponential manner throughout the whole of the Triblade under test. Uniform exponential decay would be exact if the Triblade were composed of a uniform material. Even given that the Triblade is not composed of a uniform

material, this approximation can still be plausible if the attenuation caused by the Triblade is distributed somewhat evenly among the blades that compose it, which seems reasonable [29]. See [13] for details.

2) Software Test Setup: The test applications for the Cell included five computational test codes detailed in the next paragraph (hybrid Linpack, correlator, a conjugate gradient solver, VPIC, and an integer adder) and an idle test code in which the Opteron interrogates the Cell to determine if all processor elements (PPU and SPUs) are all still operational. For the Opterons the test applications included an Opteron-only version of the correlator test code, idling, and running the Linux top command, which is considered an idle condition in the analyses that follow.

Hybrid Linpack performs the Linpack benchmark calculation, optimized for the Triblade architecture with most of the computation performed on the Cell [30], [31]. The correlator test code performs a multiply and accumulate needed for certain radio-astronomy applications [32]. It utilizes both the Opteron and PPU in very limited ways, with most of the computation performed on the SPEs. The Opteron-only correlator test code performs the same multiply and accumulate on a single Opteron core, with both cores running the code during the testing. The conjugate gradient method is a member of a family of iterative solvers used primarily on large sparse linear systems arising from the discretization of partial differential equations. The conjugate gradient test code used here performs a double precision, preconditioned conjugate gradient (CG) algorithm and utilizes the Opteron primarily for generation of the sparse linear system, with the CG implementation taking place on the Cell. VPIC is a 3D electromagnetic relativistic particle-in-cell plasma physics simulation code [33]. The version used for this testing was written to run on the Cell microprocessor in a hybrid microprocessor environment like that of a Triblade. The integer add test code is a simple hybrid code that executes primarily on the SPUs, using vector integer units to perform simple adds. Vector registers on the SPUs are loaded, vector adds are executed over these registers and verified for correctness.

Each test code was designed so that it completed its work in roughly one minute. The software setup was instrumented to run the test code continuously and return output data each time the test code completed. The output data included start and stop times, the application being run, the hardware running it, whether an SDC occurred and with the Cell idle code whether the Cells under test were still responding.

It should be noted that initial testing of the devices to determine the static sensitivity of the caches and registers to SEU was not undertaken. All of the cross-sections determined from this study are based on the dynamic usage of the system.

TABLE I
HARDWARE NEUTRON EXPOSURE

Blade	Beam Aim	Corrected Neutron Fluence ($\frac{neutrons}{cm^2}$)
Triblade 1 LS21	Upper Opteron	1.12×10^8
Triblade 1 LS21	Lower Opteron	4.20×10^8
Triblade 1 QS22a	Upper Cell	1.85×10^9
Triblade 1 QS22a	Lower Cell	1.50×10^9
Triblade 1 QS22b	Upper Cell	2.30×10^9
Triblade 1 QS22b	Lower Cell	1.87×10^9
Triblade 2 LS21	Upper Opteron	0
Triblade 2 LS21	Lower Opteron	0
Triblade 2 QS22a	Upper Cell	0
Triblade 2 QS22a	Lower Cell	1.54×10^9
Triblade 2 QS22b	Upper Cell	0
Triblade 2 QS22b	Lower Cell	1.91×10^9
Triblade 3 LS21	Upper Opteron	0
Triblade 3 LS21	Lower Opteron	0
Triblade 3 QS22a	Upper Cell	1.44×10^{10}
Triblade 3 QS22a	Lower Cell	3.51×10^9
Triblade 3 QS22b	Upper Cell	1.79×10^{10}
Triblade 3 QS22b	Lower Cell	4.36×10^9
Triblade 4 LS21	Upper Opteron	0
Triblade 4 LS21	Lower Opteron	5.37×10^8
Triblade 4 QS22a	Upper Cell	5.51×10^9
Triblade 4 QS22a	Lower Cell	2.69×10^{10}
Triblade 4 QS22b	Upper Cell	6.86×10^9
Triblade 4 QS22b	Lower Cell	3.35×10^{10}

B. Experimental Procedure

For a given experiment, a single Cell or Opteron was configured to run the desired application while the beam was aimed so that it irradiated all of the hardware in that microprocessor's beampath. With two QS22s in a Triblade, when a Cell in one QS22 is running an application, the corresponding Cell in the other QS22 is in the beampath. This second Cell in the beampath was set to run the Cell idle test code. Since the beam irradiated a columnar volume within the Triblade under test and the BC-H, certain attribution of an error to the Cells or Opteron in the beampath is not possible. In particular, other hardware in the beampath or hardware

that was affected by scatter could be the cause of an observed error. Errors could also be the result of causes external to the beam.

The experimental protocol was to start the appropriate test application on the appropriate microprocessor while the beam was off. Once the test application was observed to be operating properly (e.g., the test code had produced one or more output lines), the beam was started. The experiment continued until a state of system inoperability (e.g., a system or application crash) was reached or until sufficient time had elapsed. The beam was then turned off, data pertaining to neutron fluence exposure were collected, and the system was rebooted before beginning the next test.

For the Cells, the test procedure was to cycle through the test applications on a particular Cell, typically until it became inoperable. Repeating each test code periodically permits investigation of any aging or dose-related effects related to increasing exposure to the beam. The procedure for the Opterons, which received much less testing, was to use both the Opteron-only correlator code and possibly an idle condition (idling or running the Linux top command). Functionality of the Opteron while it was idling or running the Linux top command was assessed by ascertaining its continued responsiveness.

In all, 112 experiments were performed, with 14 Cells and 3 Opterons in four Triblades operated in the beam. The first three experiments, which were the only data collected for Triblade 2, were omitted from the results since these tests had three Triblades in the beam whereas the remaining experiments had only a single Triblade in the beam. The Opteron beampath tests were performed after the Cell beampath tests since the Cells were of primary interest in the testing. Thus, the behavior of the Opteron and the hardware in its beampath without previous exposure to the beam cannot be estimated based on this testing.

Two different beam diameters were used for the experiments: a two-inch beam diameter for the first 53 experiments and a one-inch beam diameter for the remaining 59. All testing was performed at nominal voltages and nominal temperatures with the test fixture at normal incidence to the beam. Table I details the corrected neutron fluence accumulated at each beam aim during the testing for neutrons with energies greater than 10 MeV. As it was not possible to test one Cell in a Triblade without exposing a second Cell in the beampath, the fluences include the exposure gained when a Cell was running the idle test code while the other Cell in its beampath was under test. A small amount of the Opteron memory was in the beam when the Cells were being tested, with more exposure resulting when using the two-inch beam diameter as opposed to the one-inch beam diameter. Using the two-inch beam diameter versus the one-inch beam diameter does not significantly change the hazard rate or instantaneous error rate (see Section IV D), suggesting that any resulting effects in the Opteron memory are not likely to be substantial.

IV. Results

A. Silent Data Corruption

Four SDCs were observed. Two SDCs occurred when a Cell was running a computational test code and two SDCs occurred when an Opteron was running the Opteron-only correlator test code. Checks that determined when an SDC occurred were performed using either a 160-bit secure hash algorithm-1 (SHA-1) hash or a 32-bit cyclic redundancy code (CRC), so the magnitude of the difference between the calculated result and the correct result for a particular SDC cannot be determined.

As mentioned earlier, with the Triblade architecture when one Cell is in the beam a second Cell is too. The cross-sections here and in the next subsection reflect the sensitivity of half the hardware in the Cell beampath, that is one Cell and half the remaining hardware in the Cell beampath, with this definition of the hardware in the Cell beampath applying only to cross-sections. Since with the beam diameters used for this study only a single Opteron is irradiated at a time, throughout the Opteron beampath refers to all hardware in the Opteron beampath. Further, all cross-sections reflect neutrons with energies over 10 MeV.

The cross-section for SDC for the Cell beampath is $3.88 \times 10^{-11} \frac{cm^2}{device}$ with 95% confidence interval ($3.88 \times 10^{-12} \frac{cm^2}{device}$, $1.40 \times 10^{-10} \frac{cm^2}{device}$). The cross-section for SDC for the Opteron beampath is $2.78 \times 10^{-9} \frac{cm^2}{device}$ with 95% confidence interval ($2.78 \times 10^{-10} \frac{cm^2}{device}$, $1.00 \times 10^{-8} \frac{cm^2}{device}$). These values indicate that SDC is 72 times more likely for the Opteron beampath than the Cell beampath. For the Opteron beampath, the median posterior probability that an error is an SDC rather than a failure is 0.114 with 95% credible interval (0.035, 0.250), while for the Cell beampath it is 0.038 with 95% credible interval (0.011, 0.088).

We suspect that the reason why the Opteron is more likely to incur SDC than the Cell reflects differences in the two architectures. The Cell architecture is fairly simple with a PowerPC in the PPE and a vector processing architecture in the SPEs. In both the PPE and the SPEs the processing is dependent on programmer and compiler optimization to speed up execution and increase instruction-level performance. The Opteron architecture leverages additional concepts from computer ar-

chitecture and microprocessor design to provide a high-performance microprocessor. To this end, the Opteron architecture is dependent on hardware modifications for increasing performance, including superscalar circuitry that provides real-time instruction reordering for increased performance, large multi-level cache structures for increased memory locality, and deep pipelining of functional units for increased throughput. While all of these structures provide high-performance computation, it is possible that the extra memory and logic in these functional units could be increasing the risk of SDC.

Thus, a possible explanation for the difference between the Cell beampath results and the Opteron beampath results is that the Cell architecture is much simpler than the Opteron architecture. In addition, the Opteron has two microprocessors on the same die, which tends to double the rate of both SDC and failures. Finally, as many current microprocessors are designed in the same manner as the Opteron, it is possible that SDC is more common than realized in traditional computing systems.

B. Failure Data

Each experiment was categorized as having one of two end states: 1) survival, meaning that the experiment ended when the experimenter believed the application was still running or 2) failure, indicating that the application was no longer running at the end of the experiment, e.g. because of an application or a system crash. Since the output from the test applications appeared roughly every minute, it is possible that in some cases in which the system is deemed to have survived the experiment it had actually failed, but that failure was not detected before the experiment ended. Post-irradiation analysis showed that 79 of the 95 tests conducted on the Cells ended in failure, while all 14 tests conducted on the Opterons ended in failure.

The cross-section for failures for the Cell beampath is $1.44 \times 10^{-9} \frac{cm^2}{device}$ with 95% confidence interval ($1.12 \times 10^{-9} \frac{cm^2}{device}$, $1.77 \times 10^{-9} \frac{cm^2}{device}$). When compared to the Cell beampath SDC cross-section, the Cell beampath failure cross-section indicates that failure (crashing) is 37 times more likely than SDC for the Cell beampath. The cross-section for failures for the Opteron beampath is $1.24 \times 10^{-8} \frac{cm^2}{device}$ with 95% confidence interval ($6.83 \times 10^{-9} \frac{cm^2}{device}$, $2.09 \times 10^{-8} \frac{cm^2}{device}$). When compared to the Opteron beampath SDC cross-section, the Opteron beampath failure cross-section indicates that failure is 4.5 times more likely than SDC. When the two failure cross-sections are compared, the values indicate that failure is 8.6 times more likely for the Opteron beampath than the Cell beampath. As discussed earlier, it is possible that these results reflect architectural differences.

C. Results from Modeling the Data

A model that accounted for the application used for each test, the Triblade under test, the beam aim (Cell beampath or Opteron beampath), and the beam diameter was fit to the experimental data [13]. The modeling results presented below pertain to the model used for the data and the conditions under which the experiments were conducted, with results likely to be obtained under other conditions less clear. All results have been estimated via Markov Chain Monte Carlo [34]. The paragraphs below discuss the effects of increasing exposure to the beam, beam aim, Triblade under test, application used for the test, and beam diameter.

The baseline hazard rate appears to be close to constant, suggesting that the instantaneous error rate, where errors include both failures and SDC, likely doesn't vary much with increasing exposure to the beam for the exposures observed in our testing. Therefore, it is likely that the sensitivity to radiation does not change with increasing dose accumulation and in-field usage should have roughly constant radiation-induced error rates.

The posterior probability that the beam aim (Cell beampath, which here includes both Cells and all hardware in the Cell beampath, or Opteron beampath) affects the hazard rate is 1.0. With the Opteron beampath, the median multiplier to the hazard rate is 8.785 with 95% credible interval (CI) of (4.234, 17.103), meaning there is roughly an order of magnitude more risk of an error when testing the Opteron beampath versus the Cell beampath. As previously discussed, it is possible that this disparity results from differences between the more complex architecture of the Opteron microprocessor and the simpler vector architecture of the Cell microprocessor.

There is a relationship between the Triblade under test and the beam diameter used for the testing. Triblade 3 was tested using the two-inch beam diameter and Triblade 4 was tested using the one-inch beam diameter, while Triblade 1 was tested using both beam diameters. With a situation like this, it can be difficult for the model to determine which of Triblade under test or beam diameter is more influential on the hazard rate. That said, the posterior probability that one or both of Triblade under test and beam diameter affects the hazard rate is 0.96, and the results below suggest that Triblade under test is more likely to affect the hazard rate.

The modeling results indicate a 0.880 posterior probability that different Triblades under test experienced different sensitivities to the beam. The posterior median relative difference in hazard rate for two randomly-selected Triblades is 1.335 with 95% credible interval (1.000, 4.721). Thus, this test data suggests that process-

variation-based differences in radiation sensitivity exist. However, more Triblades would need to be tested and/or more time spent under test would be required to fully explore the implications of process-variation-based radiation sensitivities.

Beam diameter (one-inch versus two-inch), on the other hand, had a 0.332 posterior probability of affecting the hazard rate, suggesting that beam diameter did not have much if any impact on the hazard rate. This implies that most of the sensitive hardware likely lies within the one-inch beam diameter.

For the most part, the application being run did not affect the hazard rate. The largest effect on the hazard rate is for hybrid Linpack, with a 0.658 posterior probability of having a different hazard rate compared to the idle condition. Its median multiplicative effect on the hazard rate is 1.434, with 95% credible interval (1.000, 2.853). Therefore, the failure sensitivity did not have much application dependence. There are a number of possible explanations for this result. First, the operating system, which executed in all tests whether an application was executing or not, might be overshadowing the effect of the application on the hardware sensitivity to neutrons. As discussed in [20] the Windows operating system increased the sensitivity to radiation-induced crashes by two to three orders of magnitude. It is possible that the Linux operating system contributes to the error rate so much that the opportunity to distinguish large application differences is not possible. Second, the applications chosen here may have similar radiation sensitivities, which other applications might not share. It would be useful to study more applications with different programming and computing patterns.

V. Conclusions

In conclusion, replicates of two microprocessors, the IBM PowerXCell 8i and the AMD Opteron 2210 HE, along with the hardware in their respective beampaths were tested at LANSCE for neutron sensitivities. These tests indicated that both microprocessor beampaths were susceptible to radiation-induced failures and SDC. The SDC cross-section was 72 times larger for the Opteron beampath than the Cell beampath and the Opteron beampath failure cross-section was almost an order of magnitude larger than that for the Cell beampath. In both cases, it is possible that these results reflect a difference in the complexity in the two architectures. The data further provided some evidence for process-variation-based radiation sensitivity differences. Finally, little application-based dependence in radiation sensitivity was found, with hybrid Linpack most likely to lead to a somewhat elevated hazard rate.

VI. Acknowledgment

Many people contributed to the success of this project. The authors thank Joe Abeyta, Chuck Alexander, Ben Bergen, Ann Borrett, Henry Brandt, James Campa, Randy Cardon, Nathan DeBardeleben, Tom Fairbanks, Parks Fields, Alan Gibson, Gary Grider, Josip Loncaric, Pablo Lujan, Alex Malin, Fred Marshall, Andrew Montoya, John Morrison, Andrew Shewmaker, Manuel Vigil, Bob Villa, Steve Wender, and Cornell Wright and apologize for any inadvertent omissions from this list.

References

[1] L. Burcin, "RAD750 experience: The challenges of SEE hardening a high performance commercial processor," in *Slides from 2002 Microelectronics Reliability and QualificationWorkshop, http://www.aero.org/conferences/mrqw/2002-papers/A Burcin*.

[2] BAE Systems. [Online]. Available: http://www.baesystems.com/

[3] J. Rhea, "BAE systems moves into third generation rad-hard processors. military and aerospace electronics." [Online]. Available: http://www.militaryaerospace.com/index/display/article-display/143464/articles/military-aerospace-electronics/volume-13/issue-5/news/bae-systems-moves-into-third-generation-rad-hard-processors.html

[4] M. Malone, "Opera RHBD multi-core," in *Military and Aerospace Programmable Logic Devices 2009*.

[5] ClearSpeed Technology plc., "Clearspeed – applications – remote and harsh environments," on web at http://www.clearspeed.com/applications/harsh/index.php.

[6] J. Ziegler and W. Lanford, "The effect of sea-level cosmic rays on electronic devices," *Journal of Applied Physics*, vol. 52, pp. 4305–4312, 1981.

[7] R. Baumann, "Radiation-induced soft errors in advanced semiconductor technologies," *IEEE Transactions on Device and Materials Reliability*, vol. 5, pp. 305–316, 2005.

[8] S. Michalak, K. Harris, N. Hengartner, B. Takala, and S. Wender, "Predicting the number of fatal soft errors in Los Alamos National Laboratory's ASC Q supercomputer," *IEEE Transactions on Device and Materials Reliability*, vol. 5, no. 3, 2005.

[9] T. Hong, S. Michalak, T. Graves, J. Ackaret, and S. Rao, "Neutron beam irradiation study of workload dependence of SER in a microprocessor," *2009 SELSE Proceedings*, 2009. [Online]. Available: http://www.selse.org

[10] K. Koch, "Roadrunner platform overview," 2008. [Online]. Available: http://www.lanl.gov/orgs/hpc/roadrunner/pdfs/Koch%20-%20Roadrunner%20Overview/RR%20Seminar%20-%20System%20Overview.pdf

[11] H. Meuer, "31st TOP500 List topped by first-ever petaflop/s supercomputer," *Scientific Computing*, 2008. [Online]. Available: http://www.scientificcomputing.com/31st-TOP500-List-Topped-by-First-ever-Petaflops-Supercomputer.aspx

[12] S. Michalak, A. DuBois, C. Storlie, W. Rust, D. DuBois, D. Modl, H. Quinn, A. Manuzzato, and S. Blanchard, "Neutron beam testing of Tribibtex NSREC-071911blades," *2011 SELSE Proceedings*, 2011. [Online]. Available: http://www.selse.org

[13] C. Storlie, S. Michalak, H. Quinn, A. DuBois, and D. DuBois, "A Bayesian survival analysis of neutron-induced errors in high performance computing hardware," Los Alamos National Laboratory, Tech. Rep., 2011.

[14] D. Hiemstra, A. Baril, and M. Dettwiler, "Single event upset characterization of the Pentium (R) MMX and Pentium (R) II microprocessors using proton irradiation," *IEEE Transactions on Nuclear Science*, vol. 46, no. 6, pp. 1453–1460, December 2009.

[15] D. Hiemstra and A. Baril, "Single event upset characterization of the Pentium (R) MMX and Celeron (R) microprocessors using proton irradiation," in *IEEE Radiation Effects Data Workshop*, 2000, pp. 39–44.

[16] D. Hiemstra, S. Yu, and M. Pop, "Single event upset characterization of the Pentium (R) MMX and Low Power Pentium (R) MMX microprocessors using proton irradiation," in *IEEE Transactions on Nuclear Science*, 2001, pp. 32–37.

[17] G. Swift, F. Farmanesh, F. Guertin, F. Irom, and D. Millward, "Single-event upset in the PowerPC750 microprocessor," *IEEE Transactions on Nuclear Science*, vol. 48, no. 6, pp. 1822–1827, December 2001.

[18] M. Cabanas-Holmen, E. Cannon, A. Kleinosowski, J. Ballast, J. Killens, and J. Socha, "Clock and reset transients in a 90 nm RHBD single-core Tilera processor," *IEEE Transactions on Nuclear Science*, vol. 56, no. 6, pp. 3505–3510, 2009.

[19] X. Li and J.-L. Gaudiot, "Tolerating radiation-induced transient faults in modern processors," *International Journal of Parallel Programming*, vol. 38, no. 2, pp. 85–116, 2010.

[20] F. Irom, "Guideline for ground radiation testing of microprocessors in the space radiation environment," Jet Propulsion Laboratory, Tech. Rep. 13, 2008.

[21] Z. Kalbarczyk, R. Iyer, and L. Wang, "Application fault tolerance with armor middleware," *IEEE Internet Computing*, pp. 28–37, March/April 2005.

[22] D. S. Katz, "Application-based fault tolerance for spaceborne applications," on Web http://trs-new.jpl.nasa.gov/dspace/bitstream/2014/10574/1/02-2631.pdf.

[23] M. Gschwind, P. Hofstee, B. Flachs, M. Hopkins, Y. Watanabe, and T. Yamazaki, "A novel simd architecture for the cell heterogeneous chip-multiprocessor," 2005, [Online]. [Online]. Available: http://www.hotchips.org/archives/hc17/2_Mon/HC17.S1/HC17.S1T1.pdf

[24] AMD, "The amd opteron processor for high-end embedded designs," 2008, [Online]. [Online]. Available: http://www.amd.com/es/Documents/42350B_Opteron_PB.pdf

[25] S. Michalak, A. DuBois, C. Storlie, W. Rust, D. DuBois, D. Modl, H. Quinn, A. Manuzzato, and S. Blanchard, "Neutron beam testing of triblades," in *Silicon Errors in Logic - System Effects 2011*.

[26] IBM. [Online]. Available: on web: http://publib.boulder.ibm.com/

[27] on web: http://www.xilinx.com, last accessed on July 11, 2011.

[28] H. Quinn, P. Graham, M. Wirthlin, B. Pratt, K. Morgan, M. Caffrey, and J. Krone, "A test methodology for determining space readiness of Xilinx SRAM-based FPGA devices and designs," *IEEE Transactions on Instrumentation and Measurement*, vol. 58, no. 10, pp. 3380–3395, 2009.

[29] S. Wender, personal communication, 2010.

[30] M. Kistler, J. A. Gunnels, D. A. Brokenshire, and B. Benton, "Programming the Linpack benchmark for Roadrunner," *IBM Journal of Research and Development*, vol. 53, pp. 9–11, 2009.

[31] A. Petitet, R. C. Whaley, J. Dongarra, and A. Cleary, "Hpl - a portable implementation of the high-performance linpack benchmark for distributed-memory computers." [Online]. Available: http://netlib.org/benchmark/hpl/

[32] A. DuBois, C. Conner, S. Michalak, G. Taylor, and D. DuBois, "Application of the IBM Cell processor to real-time cross-correlation of a large antenna array radio telescope," Los Alamos National Laboratory, Tech. Rep., 2009.

[33] K. J. Bowers, B. J. Albright, L. Yin, W. Daughton, V. Roytershteyn, B. Bergen, and T. J. T. Kwan, "Advances in petascale kinetic plasma simulation with VPIC and Roadrunner," *Journal of Physics: Conference Series*, vol. 180, no. 1, p. 012055, 2009.

[34] A. Gelman, J. Carlin, H. Stern, and D. Rubin, *Bayesian Data Analysis*, 2nd ed. Chapman & Hall, London, 2004.